

The role of top-down suppression in mitigating the disruptive effects of task-irrelevant feature changes in visual working memory

Andrea Bocincova¹ · Amanda E. van Lamsweerde¹ · Jeffrey S. Johnson¹

Published online: 24 July 2017
© Psychonomic Society, Inc. 2017

Abstract Studies of change detection have shown that changing the task-irrelevant features of remembered objects impairs change detection for task-relevant features, a phenomenon known as the irrelevant change effect. Although this effect is pronounced at short study-test intervals, it is eliminated at longer delays. This has prompted the proposal that although all features of attended objects are initially stored together in visual working memory (VWM), top-down control can be used to suppress task-irrelevant features over time. The present study reports the results of three experiments aimed at testing the top-down suppression hypothesis. Experiments 1 and 2 tested whether the magnitude or time course of the irrelevant change effect was affected by the concurrent performance of a demanding executive load task (counting backwards by threes). Contrary to the top-down suppression view, the decreased availability of executive resources did not prolong the duration of the irrelevant change effect in either experiment, as would be expected if these resources were

necessary to actively suppress task-irrelevant features. Experiment 3 showed that a visual pattern mask eliminates the irrelevant change effect and suggests that the source of the effect may lie in the use of a high-resolution, sensory memory representation to match the memory and test displays when no task-irrelevant feature changes are present. These results suggest that the dissipation of the irrelevant change effect over time likely does not depend on the use of top-down control and raises questions about what can be inferred about the nature of storage in VWM from studies of this effect.

Keywords Feature binding · Irrelevant change effect · Visual working memory · Top-down suppression · Change detection task

Introduction

Working memory is a severely capacity-limited system that plays a critical role in many daily tasks, including searching for objects (Carlisle, Arita, Pardo, & Woodman, 2011), language processing (Gregoriou, Gotts, Zhou, & Desimone, 2009), mental calculations (Logie, Gilhooly, & Wynn, 1994), and retrieving information from long-term memory (Unsworth & Engle, 2007). Two outstanding and interrelated issues that are of central importance to our understanding of this important cognitive system include the format of the visual working memory (VWM) representation (i.e., what does a representation “look like?”) and the extent to which top-down control can be exerted over what is encoded and maintained in VWM. One behavioral effect, the *irrelevant change effect*, has been proposed to shed light on both of these issues. The irrelevant change effect occurs when detecting changes to task-relevant features is impaired due to changes in task-irrelevant features. This effect is observed in the context of

Andrea Bocincova and Amanda E. van Lamsweerde contributed equally to this article.

Electronic supplementary material The online version of this article (doi:10.3758/s13421-017-0738-8) contains supplementary material, which is available to authorized users.

✉ Amanda E. van Lamsweerde
amanda.vanlamsweerde@nds.u.edu

Andrea Bocincova
andrea.bocincova@nds.u.edu

Jeffrey S. Johnson
jeffrey.s.johnson@nds.u.edu

¹ Department of Psychology and Center for Visual and Cognitive Neuroscience, North Dakota State University, Department 2765, P.O. Box 6050, Fargo, ND 58108, USA

change detection tasks in which participants are asked to remember a small set of multi-featured objects (e.g., colored shapes) over a brief delay, and then report whether items in a test display are the same as or different to the items they remembered viewing in the initial display. In the irrelevant-change version of this task, participants are instructed to remember and detect changes to a single attribute of each stimulus (e.g., its color) while ignoring other stimulus attributes (e.g., shape). On some trials, however, changes may also occur along the task-irrelevant dimension (e.g., the original shape could be replaced by a new shape). The logic is that if the stored representation includes both task-relevant and task-irrelevant features of the remembered objects, performance should be disrupted when a task-irrelevant change occurs at test (the irrelevant change effect). This effect has been observed in numerous experiments to date (Ecker, Maybery, & Zimmer, 2013; Gao, Li, Yin, & Shen, 2010; Hyun, Woodman, Vogel, Hollingworth, & Luck, 2009; Logie, Brockmole, & Jaswal, 2011; Shen, Tang, Wu, Shui, & Gao, 2013; Treisman & Zhang, 2006; Woodman, Vogel, & Luck, 2012; Yin et al., 2011; Yin, Gao, et al. 2012; Zhou et al., 2011), and has been taken as evidence for the proposal that the separate features of attended objects, whether task-relevant or not, are automatically bound together and stored in VWM as integrated object-based representations. Thus, the standard interpretation of these results is that the irrelevant change effect reflects a disruption of performance caused by changing task-irrelevant features at test, which occurs because VWM is object-based.

Although the irrelevant change effect has been observed in numerous cases, research by Jaswal and colleagues (Jaswal & Logie, 2011; Logie et al., 2011; see also Treisman & Zhang, 2006) demonstrated that the presence of the effect depends on the length of the delay between the offset of the initial memory stimulus and the appearance of the test display (i.e., on the interstimulus interval, or ISI, between the memory and test displays). Specifically, the irrelevant change effect was observed when the memory-test ISI was short (~0–1,000 ms) but not when it was long (>~1,500 ms). The authors explained the time dependency of this effect by proposing that the spatial and surface features of objects are initially bound together via low-level perceptual processes, and stored as object files in VWM; however, over time, top-down control processes can be used to selectively inhibit object properties that are not relevant to current task performance, and whose inclusion in the memory representation may prove disruptive (Jaswal, 2012). The purpose of the present study was to directly test whether the elimination of the irrelevant change effect over time depends on the availability of executive resources that serve to suppress task-irrelevant features. More specifically, we were interested in testing whether the disruptive effects of irrelevant feature changes can be regulated by the use of domain-general executive resources, which are presumed to

form a part of the broader working memory system (see, e.g., Baddeley, 2012). Executive resources of this kind have been proposed to play a general role in the mental manipulation of information in working memory, in addition to supporting the suppression of distracting information in other memory tasks (Conway & Engle, 1994; Kane & Engle, 2003). Additionally, in their concluding paragraph, Logie et al. (2011) explicitly proposed this form of executive resource as potentially mediating the inhibition of task-irrelevant features over time.

To test this hypothesis, observers were presented with two arrays of colored shapes separated by a blank delay interval that varied in duration (see Fig. 1). The observer's task was to report whether the two arrays were the same or if the binding of color and location (Experiment 1) or shape and location (Experiment 2) differed between sample and test. Across separate sessions, observers completed either a standard or an irrelevant change version of the task. In the standard version of the task, the task-irrelevant feature of each object (shape in Experiment 1 and color in Experiment 2) remained unchanged from sample to test; in the irrelevant change variant of the task, the irrelevant feature of each object was replaced with a different feature between sample and test. Additionally, across trials in each session, the primary task was paired with one of two different secondary tasks (randomly determined): either a low-load articulatory suppression task (verbally repeating a single word throughout each trial), or a high-load counting task (counting backward by threes), a commonly used means of taxing executive resources (see, e.g., Allen, Baddeley, & Hitch, 2006).

This task design made it possible to determine whether the decline in the irrelevant change effect over time depends on the availability of executive resources during the delay. According to the proposal of Logie et al. (2011), if executive resources are not available to inhibit the task-irrelevant features, changes in these features at test should continue disrupting performance even at long sample-test delays; that is, the irrelevant change effect should persist over time in the high-load condition. A key assumption of this hypothesis is that the irrelevant change effect is caused by a disruption of performance on irrelevant change trials. If the irrelevant change effect is caused by a failure to ignore irrelevant feature changes at test, as most previous research has assumed, then observed reductions in the magnitude of the effect over time should be accounted for primarily by a decrease in the likelihood of making an error on irrelevant-change trials; that is, of incorrectly responding “change” when no task-relevant change has occurred (a false-alarm response), or of incorrectly responding “same” when a task-relevant change *has* occurred (a miss response). The results of both Experiment 1 and Experiment 2 indicated that the availability of executive resources of the sort manipulated here is not crucial in order to observe time-based changes in the irrelevant change effect; therefore, Experiment 3 was designed to test the alternate

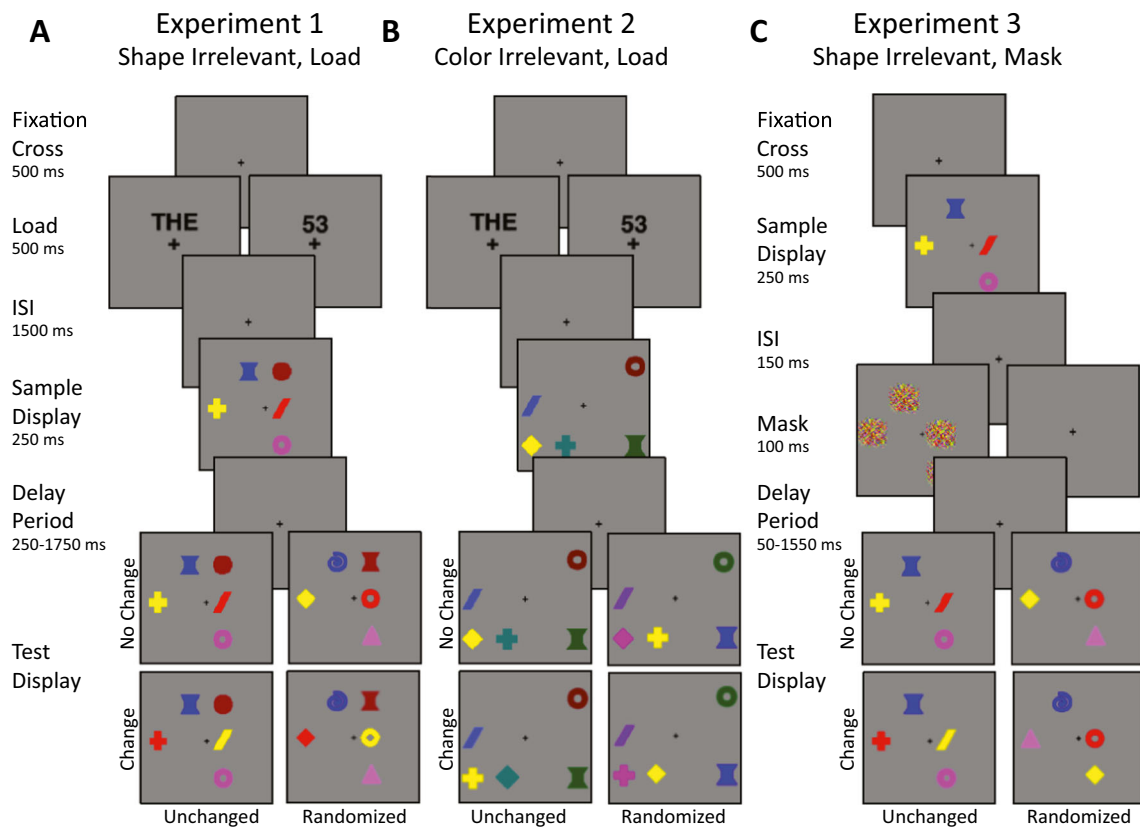


Fig. 1 Behavioral task trial sequence. Participants performed a visual working memory (VWM) change detection task in two experimental sessions. In one of the sessions, task-irrelevant feature was randomized between sample and test, i.e., shape in Experiment 1 (a) and Experiment 3

(c), and color in Experiment 2 (b). Load manipulation was included in Experiment 1 (a) and Experiment 2 (b). Mask manipulation was included in Experiment 3 (c). Stimuli are not drawn to scale

hypothesis that the irrelevant change effect is caused by a fast-matching process. This experiment is described in greater detail following Experiment 2.

Experiment 1

The purpose of Experiments 1 and 2 was to test the hypothesis that VWM automatically encodes all features of an object and that, over time, executive resources can be engaged to inhibit task-irrelevant features. The experimental design was similar to that of Logie et al. (2011), in which participants were required to remember color-location or shape-location bindings, while the third dimension was task irrelevant. The length of the delay interval was variable and ranged from 250 to 1,750 ms. Additionally, on half the trials, executive resources were occupied with a concurrent backward counting task. Participants were instructed to count backward by threes (high-load condition) or repeat the word “the” throughout the trial (low-load condition). If top-down suppression of task-irrelevant features is responsible for the decline in the irrelevant change effect over time, then the high-load task should prevent suppression and the irrelevant change effect

should be observed across all delay periods. In Experiment 1, participants were asked to remember color-location bindings and, in one of the two experimental sessions, to ignore task-irrelevant changes in shape.

Method

Participants Twenty-seven undergraduate and graduate students (mean age = 19.9 years, SD = 3.3, 15 female) from North Dakota State University participated in this experiment for course credit or monetary compensation (US\$10/h). All participants reported normal or corrected-to-normal visual acuity and normal color vision. One participant was dropped due to poor performance on both the primary and the concurrent task (performance was consistently at chance levels across experimental blocks and no more than one subtraction was performed on the majority of trials) and two participants were excluded from subsequent analysis because they failed to complete both experimental sessions, resulting in a final N of 24 participants. Each participant provided written informed consent. All experimental protocols were approved by the North Dakota State University Institutional Review Board (IRB).

Stimuli Sample displays contained five objects randomly positioned at different locations within an imaginary 3×4 grid subtending $6.1 \times 7.8^\circ$ of visual angle at a viewing distance of 60 cm. Objects were created by randomly sampling without replacement from a set of 12 shapes (swirl, star, triangle, hourglass, circle, square, cross, horseshoe, t-shape, donut, diamond, slash) and six colors with the following coordinates in the 1931 CIE color coordinate system: red ($x = 0.61, y = 0.34; 14.51 \text{ cd/m}^2$), yellow ($x = 0.41, y = 0.51; 63.41 \text{ cd/m}^2$), green ($x = 0.29, y = 0.61; 48.92 \text{ cd/m}^2$), blue ($x = 0.15, y = 0.07; 8.72 \text{ cd/m}^2$), pink ($x = 0.27, y = 0.17; 23.43 \text{ cd/m}^2$), wine ($x = 0.57, y = 0.32; 3.90 \text{ cd/m}^2$). Stimuli were presented against a gray background ($x = 0.28, y = 0.30; 21.10 \text{ cd/m}^2$) on the surface of a 21-in. CRT monitor.

Four different test displays were constructed, depending on the trial type. In the *unchanged-shape* condition, the shapes remained in the same locations from study to test. For half of these trials, the test displays were identical to the sample display (no-change trials), while on the other half two colors swapped their locations between sample and test (change trials). In the *randomized-shape* condition, the shape of each object was replaced with a different shape chosen at random, without replacement, from the set of 12 possible shapes. For half of these trials, the color-location bindings were identical to the sample display (no-change trials), while on the other half, two colors swapped their locations between sample and test (change trials). In each case, the participants' task was to detect whether the binding of color and location for any two items had changed between sample and test, irrespective of whether an irrelevant shape change had occurred or not.

Procedure The basic trial design can be seen in Fig. 1a. Each trial began with a fixation cross for 500 ms, followed by a random two-digit number (>20) or the word "THE" at the center of the screen, depending on condition. In the *low-load* condition, participants simply repeated the word "the" throughout the trial at a rate of $\sim 2\text{--}3$ repetitions/s. In the *high-load* condition, participants were required to count backward from the two-digit number by threes throughout the trial. In each case, the digit or word screen was followed by a 1,500-ms interstimulus interval and the 250-ms presentation of the sample display. The sample display was followed by a delay of 250, 750, 1,250, or 1,750 ms, randomly intermixed, and then the test display, which remained visible until the participant made a response. At test, participants indicated whether color-location bindings remained the same or differed between the sample and test displays by pressing one of two buttons on a computer keyboard using either their left or right index finger.

For trials with an executive load, participants were instructed to complete approximately one subtraction per second ($\sim 1\text{--}2$ subtractions in the shortest delay condition and $\sim 3\text{--}4$ subtractions in the longest delay condition). To ensure

compliance with the load task, a research assistant was present who tracked counting accuracy and, when necessary, prompted participants to continue counting throughout the full delay period. High- and low-load trials were randomly intermixed and occurred equally often within experimental blocks. The unchanged-shape and randomized-shape conditions were run as separate sessions that were completed on different days (minimum separation = 1 day, maximum = 5 days), with session order counterbalanced across participants. Each session lasted ~ 1 h and was comprised of 40 trials in each load \times delay condition (20 same and 20 different trials/condition; 320 trials in total), grouped into ten blocks of 32 trials each, in addition to 32 practice trials. Participants were given a short break after every block.

Results

We conducted analyses on d' and hits and false alarms separately.¹ Using the d' measure, we first sought to determine whether the executive load manipulation influenced the irrelevant change effect at longer delays. Finally, to clarify the source of the irrelevant change effect and its reduction over time, we conducted separate analyses of the hit and false-alarm rates.

Change detection performance (d') Mean change detection performance (d') across conditions can be seen in Fig. 2a. The analyses revealed that the high-load task reduced performance overall, but did not affect the time-based reduction of the irrelevant change effect. Specifically, d' values were analyzed with a three-way within-subjects ANOVA² with factors of load (low, high), irrelevant feature (unchanged, randomized), and delay (250, 750, 1,250, and 1,750 ms). This revealed a significant main effect of load, $F(1,23) = 48.181, p < .001, \eta_p^2 = .677$, caused by lower performance in the high-load than the low-load conditions. The main effect of irrelevant feature was also significant, $F(1,23) = 51.171, p < .001, \eta_p^2 = .690$. Change detection performance was significantly lower when the task-irrelevant feature (shape) was randomized between the sample and test displays. There was also a significant main effect of delay, $F(3,69) = 20.298, p < .001, \eta_p^2 = .469$, with significantly worse change detection performance at long-versus short-delay intervals. The irrelevant feature \times delay

¹ Response times (RTs) were also analyzed, although there was no indication that load changed the pattern of RTs (there was no interaction with load with any other factor in Experiments 1 and 2); therefore, RT results were not included in the manuscript. Briefly, in Experiments 1 and 2, RTs were longer with a load than without, and RTs were longer with increasing delay times, although this increase in RT at longer delays was only present for non-randomized displays. In Experiment 3, RTs were longer as delay increased, and the presence of the mask increased RTs for the unchanged displays, but decreased RTs for randomized displays.

² All follow-up pairwise comparisons were Bonferroni adjusted to alpha level .0125 (.05/4) unless otherwise stated.

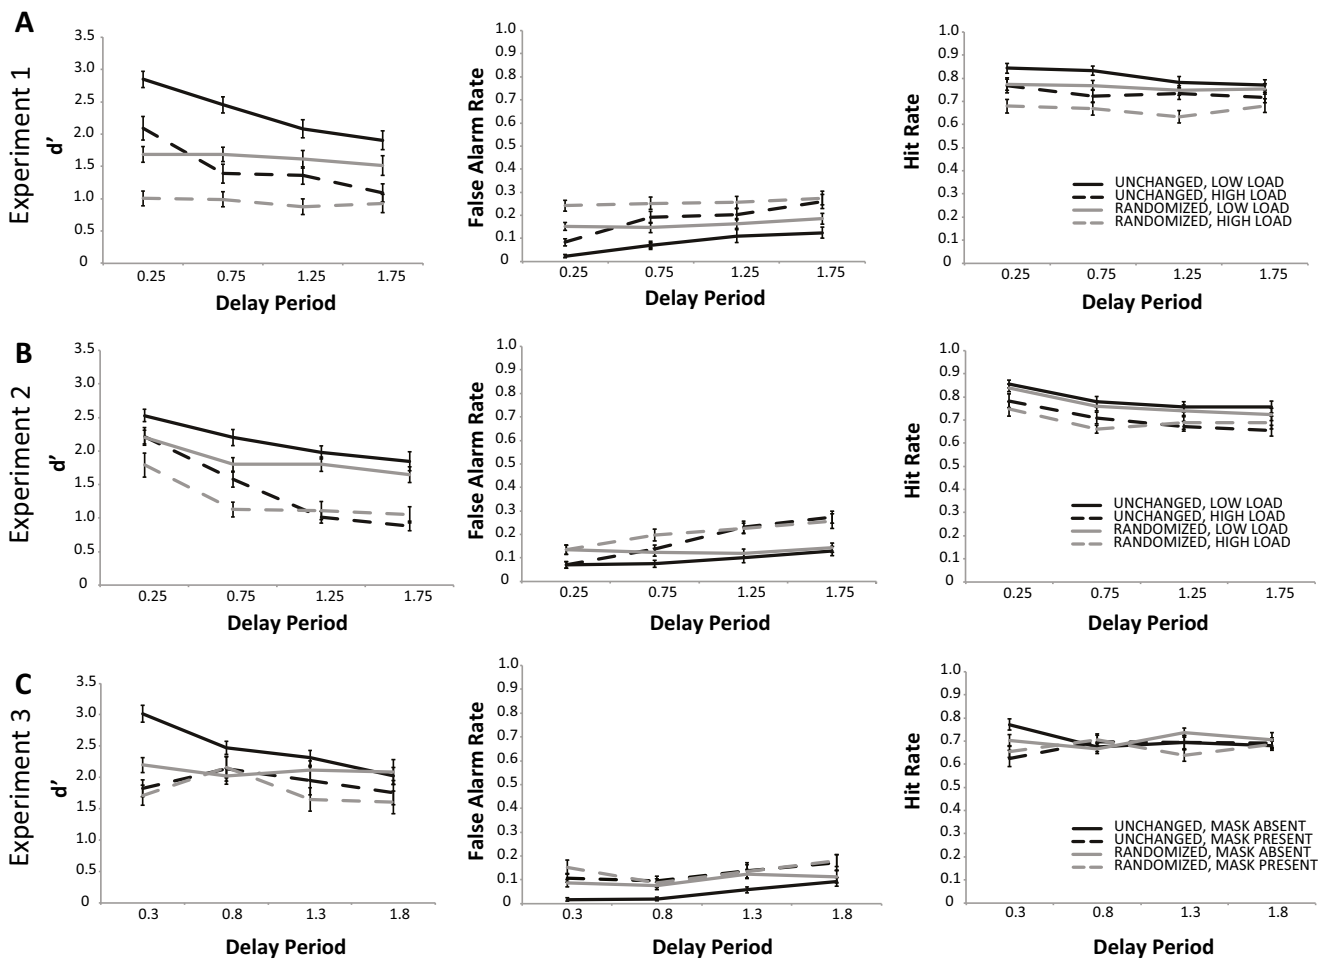


Fig. 2 Results of Experiment 1 (a), 2 (b), and 3 (c). Each column depicts a different measure of behavioral performance: d' , false-alarm rate, and hit rate, respectively. Each row corresponds to a different

experiment. Delay period lengths are expressed in seconds. Error bars represent standard error of the mean

interaction was also significant, $F(3,69) = 13.700, p < .001, \eta_p^2 = .373$. In keeping with the findings of Logie et al. (2011), tests of simple effects showed that performance (averaged across load conditions) was better in the unchanged versus randomized irrelevant-feature condition when the delay was short (pairwise comparisons for delays 250, 750, and 1,250 ms $p < .002$, Bonferroni corrected), but performance in the two conditions converged at the longest delay (pairwise comparison $p = .043$; Bonferroni corrected). Critically, none of the interactions involving the factor of load were significant (all $ps > .149$). This suggests that the irrelevant-change effect dissipated over time irrespective of whether executive resources were available or were occupied with the counting backwards task.

Hits and false alarms As noted in the introduction, changes in irrelevant features at test could affect performance in at least two different ways. First, changes in task-irrelevant features could be mistaken for task-relevant changes, inflating the

false-alarm rate in the randomized condition at short delays, before the task-irrelevant features can be removed from the memory representation (or, alternately, be allowed to passively decay). Alternately, task-irrelevant changes could make it difficult to detect task-relevant changes at test, which would lead to a reduction in hits in the randomized condition. Another possibility is that the reduction of the irrelevant change effect over time is driven by a drop in performance in the unchanged condition (i.e., by an increase in false-alarm rate or a decrease in hit rate), rather than improved performance in the randomized condition. To assess these possibilities, hits and false-alarm rates were analyzed using separate three-way within-subjects ANOVAs² with factors of load (low, high), irrelevant feature (unchanged, randomized), and delay (250, 750, 1,250, and 1,750 ms). The results suggest that the irrelevant change effect is caused by a decrease in performance across time for the unchanged displays, while performance for the randomized displays remained stable over time.

The ANOVA results for false-alarm rate revealed significant main effects of load, delay, and irrelevant feature (all p s < .001). Additionally, there was a significant irrelevant feature \times delay interaction, $F(3, 69) = 21.273$, $p < .001$, $\eta_p^2 = .262$. Tests of simple effects comparing the effect of irrelevant feature (unchanged, randomized) at each delay suggested that the interaction was driven by a significant difference in the effect of irrelevant feature on false-alarm rates at the two shortest delays (pairwise comparisons for 250 and 750 ms p s < .001). However, the effect was not driven by a reduction in false alarms in the randomized condition, as would be expected if the disruptive effect of irrelevant feature changes was reduced over time; instead, as can be seen in Fig. 2a, differences between conditions appear to be driven primarily by an extremely low false-alarm rate in the unchanged condition at the earliest delays, which increased steadily from 750 to 1,750 ms, at which point the false-alarm rates are essentially identical.

By contrast, analysis of hit rates revealed significant main effects of irrelevant feature [$F(1,23) = 14.708$, $p < .001$, $\eta_p^2 = .390$] and load [$F(1,23) = 28.700$, $p < .001$, $\eta_p^2 = .550$], but not of delay ($p = .084$). Additionally, there were no significant interactions (all p s > .190). These results suggest that although randomization of the irrelevant feature and performance of the high-load secondary task reduced the likelihood of correctly detecting a change at test, the magnitude of these effects did not differ as a function of delay.

Discussion

The goal of Experiment 1 was to test whether the reduction of the irrelevant change effect over time depends on the availability of executive resources to suppress irrelevant features in VWM. In contrast to this prediction, a resource-demanding secondary task caused a large drop in overall performance but did not prevent the irrelevant change effect from dissipating at longer delay intervals. Additionally, analyses of hits and false alarms showed that the reduction in the irrelevant change effect over time was most likely due to a decrease in correct rejections (i.e., an increase in false alarms) in the unchanged irrelevant feature condition – i.e., in the likelihood of correctly responding “same” when no task-relevant or irrelevant change occurred – rather than a decrease in false alarms (or increase in hits) in the randomized irrelevant feature condition. Further discussion of the implications of these findings will be delayed until after Experiment 2.

Experiment 2

Experiment 2 was identical to Experiment 1, with the exception that, instead of remembering color-location bindings, participants were asked to remember shape-location bindings and to ignore color.

Method

Participants Thirty undergraduate students (mean age = 19.3 years, SD = 2.2, 16 female) from North Dakota State University participated in this experiment for course credit. All participants reported normal or corrected-to-normal vision. Six participants were excluded from subsequent analysis because they failed to complete both experimental sessions, resulting in a final N of 24. Each participant provided written informed consent. All experimental protocols were approved by the North Dakota State University IRB.

Stimuli and procedure The stimuli and procedure were identical to Experiment 1, except that stimuli were randomly drawn from a set of six shapes (hourglass, cross, t-shape, donut, diamond, slash) and 12 colors: yellow ($x = 0.42$, $y = 0.50$; 56.75 cd/m²), orange ($x = 0.52$, $y = 0.42$; 25.61 cd/m²), green ($x = 0.29$, $y = 0.58$; 28.28 cd/m²), deep green ($x = 0.30$, $y = 0.58$; 4.06 cd/m²), olive ($x = 0.40$, $y = 0.51$; 7.34 cd/m²), sea green ($x = 0.21$, $y = 0.28$; 11.54 cd/m²), sky blue ($x = 0.16$, $y = 0.08$; 9.98 cd/m²), blue ($x = 0.15$, $y = 0.07$; 8.72 cd/m²), red ($x = 0.61$, $y = 0.34$; 11.19 cd/m²), wine ($x = 0.59$, $y = 0.34$; 2.77 cd/m²), fuchsia ($x = 0.28$, $y = 0.14$; 14.99 cd/m²), orchid ($x = 0.20$, $y = 0.10$; 9.89 cd/m²). On half of the trials, a task-relevant change in shape-location bindings was introduced; i.e., two shapes swapped their locations between sample and test. Additionally, in a randomized-color condition, the color of each object was replaced by a different color drawn randomly from the set of all possible colors.

Results

Change detection performance (d') The data for Experiment 2 are shown in Fig. 2b. As in Experiment 1, performance of the high-load secondary task reduced performance overall, but did not affect the time-based elimination of the irrelevant change effect. The effect of the load manipulation on the irrelevant change effect was assessed with a three-way within-subjects ANOVA² with factors of load (low, high), irrelevant feature (unchanged, randomized), and delay (250, 750, 1,250, and 1,750 ms). Once again, this analysis revealed a significant main effect of load, $F(1,23) = 66.787$, $p < .001$, $\eta_p^2 = .744$. Performance of the counting backwards task impaired change detection performance across all conditions. The main effect of delay was also significant, $F(3,69) = 59.172$, $p < .001$, $\eta_p^2 = .720$, reflecting the fact that performance steadily declined as the delay interval increased in length. There was also a significant main effect of irrelevant feature, $F(1,23) = 14.706$, $p = .001$, $\eta_p^2 = .390$. Randomization of the (task-irrelevant) color of each object from sample to test had a disruptive effect on overall performance. As in Experiment 1, there was also an irrelevant feature \times delay interaction, $F(3,69) = 5.290$ $p =$

.002, $\eta_p^2 = .187$. Pairwise comparisons revealed that the irrelevant change effect was only observed at the two shortest sample-test delays ($ps < .002$; Bonferroni corrected), with performance in the unchanged- and randomized-color conditions converging at delays greater than 500 ms ($ps > .698$; Bonferroni corrected). Additionally, unlike Experiment 1, there was a significant load \times delay interaction $F(3,69) = 6.598$, $p = .001$, $\eta_p^2 = .223$, reflecting a steeper drop in performance overall as a function of delay in the high-load compared to the low-load condition. However, this interaction did not reflect a difference in the magnitude of the irrelevant change effect between conditions at later delays: none of the remaining interactions reached significance (all $ps > .252$), including the three-way load \times irrelevant feature \times delay interaction, $F(3,69) = 1.394$, $p = .252$, $\eta_p^2 = .057$. As in Experiment 1, the disruptive effect of randomizing the irrelevant feature (color, in this case) was eliminated at longer delays, even when executive resources were consumed by the counting backwards task.

Hits and false alarms In keeping with the results of Experiment 1, a three-way within-subjects ANOVA² assessing the effect of load (low, high), irrelevant feature (unchanged, randomized), and delay (250, 750, 1,250, and 1,750 ms) on false alarms revealed significant main effects of each factor (all $ps < .018$), a significant load \times delay interaction [$F(3,69) = 17.491$, $p < .001$, $\eta_p^2 = .432$] reflecting a steeper increase in false alarms with increasing delay length in the high versus the low load condition, and a significant irrelevant feature \times delay interaction [$F(2,003,16.5449) = 3.911$, $p = .027$, $\eta_p^2 = .145$; Greenhouse-Geisser corrected]. Tests of simple effects revealed that the irrelevant feature \times delay interaction was driven by a significant difference in scores between the randomized and unchanged irrelevant feature conditions at the two shortest delays ($ps < .001$; Bonferroni corrected). Once again, as can be seen in Fig. 2b, the rate of false alarms was very low in the shortest delay condition, and increased steadily at longer delays. Mirroring the results of Experiment 1, this suggests that the reduction of the irrelevant change effect over time was likely caused by a drop in performance in the unchanged irrelevant feature condition – i.e., an impaired ability to correctly respond “same” when no change has occurred – rather than an improvement in the ability to ignore task-irrelevant changes in the randomized condition.

Analysis of hit rate revealed a significant main effect of load [$F(1,23) = 40.370$, $p < .001$, $\eta_p^2 = .639$] and delay [$F(1,23) = 2.307$, $p < .001$, $\eta_p^2 = .446$]. There were no significant interactions ($ps > .276$; Bonferroni corrected), suggesting that the effect of irrelevant changes on hit rate did not vary either as a function of load or of elapsed time since memory display offset.

Discussion

In keeping with previous findings and the results of Experiment 1, Experiment 2 showed that irrelevant color changes disrupted performance at short but not long study-test intervals. However, similar to the findings of Logie et al. (2011), the disruption produced at short delays was smaller than that observed for irrelevant shape changes in Experiment 1. This could be due to differences in the relative importance of color versus shape for initial binding, as proposed by Logie et al., or may simply be a consequence of the specific colors and shapes used in this particular experiment. Determining which of these possibilities best explains the data is beyond the scope of the present study. Importantly, however, in both experiments the same pattern of results was observed whether participants performed a simple articulatory suppression task, repeating the word “the” over and over throughout the primary task, or a more challenging backwards counting task concurrent with primary task performance. Although performance of the high-load counting task negatively affected performance overall, it did not prevent the irrelevant change effect from dissipating at longer delays. Finally, analysis of false alarms and hits suggests that, irrespective of load, the effect was most likely driven by a change over time in participants’ ability to correctly respond “same” when there was no change at test (either task-relevant or irrelevant). That is, at short delays and in the absence of irrelevant changes, participants were very unlikely to incorrectly respond “change” when no change had occurred (i.e., to make a false alarm).

The results of Experiments 1 and 2 revealed that the irrelevant change effect declined over time in both the low- and high-load conditions. Taken together, these results suggest that the elimination of the irrelevant change effect over time likely does not depend on the use of executive resources to suppress task-irrelevant features. This finding is inconsistent with the top-down suppression hypothesis as proposed by Logie et al. (2011). However, an alternative possibility is that task-irrelevant features decay passively over time when resources are unavailable for their storage. Support for this possibility comes from a study by Xu (2010) that used MRI to examine the obligatory encoding of task-irrelevant object features in two regions known to contribute to storage in VWM: the superior intraparietal sulcus (IPS) and the lateral occipital complex (LOC). In this experiment, participants were presented with a variable number of unique shape/color pairs, but only the color of each item was task-relevant and needed to be stored in VWM. Analysis of BOLD signal changes during the memory delay in each area revealed that the superior IPS was only sensitive to the number of task-relevant features that were maintained, but the signal in LOC was sensitive to both the number of unique colors and unique shapes, even though shape was not relevant to the task. However, when working memory load was high, BOLD signal changes reflecting the

encoding of task-irrelevant features in the LOC rapidly decayed following stimulus offset. This finding was taken as support for the proposal that, although initial encoding of task-irrelevant features may be automatic, maintenance of this information is under voluntary control. That is, active storage requires the allocation of memory resources; if these resources are occupied, task-irrelevant information is represented only transiently. This finding raises the possibility that the removal of irrelevant object properties does not require active inhibition, but may occur automatically in the absence of sufficient resources (for behavioral evidence challenging this view, see Shen et al., 2013; Yin, Zhou, et al., 2012).

However, contrary to either the active or passive suppression views, elimination of the irrelevant change effect over time was largely accounted for by a drop in performance on standard trials, rather than a reduction in errors on irrelevant-change trials. Specifically, the likelihood of making a correct “same” response in the absence of task-irrelevant changes decreased substantially at longer sample-test delays. This finding suggests that the presence of irrelevant changes at test may impair performance by preventing observers from utilizing a high-capacity but fast-decaying memory trace to efficiently match the sample and test displays when the delay is short. Specifically, when the sample-test delay is short and task-irrelevant features remain constant from sample to test, participants appear to be taking advantage of a high capacity but quickly fading representation to efficiently match the sample and test arrays. This leads to extremely accurate performance on no-change trials when the delay interval is short, and a gradual increase in false alarms as the delay grows longer and performance comes to rely on a more abstract form of visual memory. Therefore, we propose that the irrelevant change effect is likely driven by a boost in performance at early time intervals when task-irrelevant features match from study to test. To test this possibility more directly, we conducted a third experiment in which visual masks were interposed between the sample and test displays. If differences between the standard and irrelevant change variants of the task are primarily driven by the use of a highly efficient matching process in the standard task at short delays, we expected the effect to go away when this process was interrupted by the presentation of a visual mask.

Experiment 3

To test the enhanced matching hypothesis, in Experiment 3 we replicated the basic methods used in Experiments 1 and 2, but visual pattern masks were presented at the location of each remembered object during the interval between the sample and test displays. Previous research suggests that pattern masks disrupt short-lived, high-capacity forms of visual memory, such as iconic memory (Phillips, 1974) and “fragile”

VWM (Sligte, Scholte, & Lamme, 2008), but have no effect on longer lived, more capacity-limited forms of VWM (Vogel, Woodman, & Luck, 2006). Specifically, there appears to be a period of time outside of typical iconic memory, prior to consolidation of a fully durable working memory representation, where memory representations have qualities that are similar to a sensory representation: capacity is higher than working memory and the memory trace is maskable (Sligte, 2010; Sligte et al., 2008; Vandenbroucke & Sligte, 2011). While this has been conceptualized by some as a discrete stage of memory (Sligte et al., 2008), this point of view is not without criticism (Matsukura & Hollingworth, 2011). However, regardless of whether this transition between iconic and working memory is characterized as a discrete stage or a point along a continuum from a fleeting sensory representation to a durable VWM representation (Treisman & Zhang, 2006), there does appear to be a sensory trace that can be used in change detection, even outside of the traditional time limits of iconic memory (Bradley & Pearson, 2012; Sligte, 2010). If improved performance in the unchanged irrelevant feature condition reflects the use of a high-capacity short-lived form of visual memory to efficiently match the sample and test displays, we expected the presentation of pattern masks to eliminate the irrelevant change effect at all delays. By contrast, if the effect depends on the gradual removal of task-irrelevant features from VWM representations, we expected the effect to persist at short delays, and to decline gradually over time, as observed in Experiments 1 and 2.

Method

Participants Eighteen undergraduate and graduate students (mean age = 20.05 years, SD = 2.1, 8 female) from North Dakota State University participated in this experiment for course credit. Our choice of sample size was based on a post-hoc power analysis of the results of Experiment 1, which suggested that six participants would give us 80% power to detect an irrelevant change effect of similar magnitude to that observed. However, because we were predicting a null effect in the masked condition, we elected to triple this number to 18. All participants reported normal or corrected-to-normal vision. Each participant provided written informed consent. All experimental protocols were approved by the North Dakota State University IRB.

Stimuli and procedure The stimuli and procedure were identical to Experiment 1, except that the set size was decreased to four objects and a visual masking display was added on half of the trials. The masking display consisted of four patches of colored, oriented lines created using all of the possible experimental colors and presented at the locations of each of the objects comprising the sample display. The mask display was presented 150 ms after the offset of the sample display and

remained present for 100 ms. Following the mask display, the test display appeared following a variable delay (50, 550, 1,050, or 1,550 ms). Due to the presence of the mask, the ISI between the offset of the memory display and the appearance of the test display was increased by 50 ms compared to Experiments 1 and 2 (300, 800, 1,300, and 1,800 ms).

Results

The results of Experiment 3 are shown in Fig. 2c. The data reveal that masks eliminated the irrelevant change effect at all intervals, while also reducing performance overall.

Change detection performance (d') Performance (d') was analyzed using a three-way within-subjects ANOVA² with factors of mask (present, absent), irrelevant feature (unchanged, randomized), and delay (300, 800, 1,300, and 1,800 ms). The analysis revealed a main effect of mask, $F(1,17) = 26.504, p < .001, \eta_p^2 = .609$. Presentation of a mask following the sample display had an overall detrimental effect on performance. There was also a main effect of irrelevant feature, $F(1,17) = 5.702, p = .029, \eta_p^2 = .251$; performance was worse overall when task-irrelevant features were randomized between sample and test. The main effect of delay was also significant, $F(3,51) = 4.816, p = .005, \eta_p^2 = .221$, reflecting the fact that performance declined as the delay interval grew longer. Additionally, there were significant two-way interactions between mask and delay [$F(3,51) = 6.624, p = .001, \eta_p^2 = .280$], and irrelevant feature and delay [$F(3,51) = 3.766, p = .016, \eta_p^2 = .181$]. Pairwise comparisons following up on the significant mask \times delay interaction revealed significant differences in performance between the mask present versus absent conditions at three delay period lengths (300, 800, and 1,800 ms, $ps < .008$, Bonferroni corrected). Follow-up pairwise comparisons for the irrelevant feature \times delay interaction revealed an irrelevant change effect – i.e., a significant difference between randomized and unchanged conditions at the shortest delay ($p = .002$, Bonferroni corrected). Most importantly, however, there was also a significant irrelevant feature \times mask \times delay interaction, $F(3,51) = 3.174, p = .032, \eta_p^2 = .157$, suggesting that the irrelevant change effect was only present at the two shortest delays (300 and 800 ms) when no mask was presented, $t(17) = 4.203, p = .001$ and $t(17) = 3.817, p = .001$, respectively ($ps > .270$ for all other delays). No irrelevant change effect was present when visual masks were presented in the time period between the sample and test displays ($ps > .100$, Bonferroni corrected). In line with our hypothesis, the inability of participants to efficiently match the sample and test displays in the masked version of the unchanged condition eliminated the performance advantage in this condition and removed the irrelevant change effect.

Hits and false alarms In order to examine the effect of visual masks on the ability of participants to correctly identify task-relevant changes, false alarms and hits were analyzed using separate three-way within-subjects ANOVAs² with factors of mask (mask present, mask absent), irrelevant feature (unchanged, randomized), and delay (300, 800, 1,300, and 1,800 ms). Analysis of false-alarm rates revealed main effects of irrelevant feature [$F(1,17) = 6.030, p = .025, \eta_p^2 = .262$], mask [$F(1,17) = 16.010, p = .001, \eta_p^2 = .485$], and delay [$F(3,51) = 7.654, p < .001, \eta_p^2 = .310$]. Aside from the main effects, there was a significant irrelevant feature \times mask interaction [$F(1,17) = 5.415, p = .033, \eta_p^2 = .242$] suggesting that participants committed fewer false alarms when no mask was presented during the delay period in both randomized and unchanged conditions ($ps < .020$, Bonferroni corrected to alpha level .025).

Analysis of hit rates revealed a significant main effect of mask [$F(1,17) = 7.928, p = .012, \eta_p^2 = .318$], a significant two-way mask \times delay interaction [$F(3,51) = 5.213, p = .003, \eta_p^2 = .235$], and a three-way irrelevant feature \times mask \times delay interaction [$F(3,51) = 4.192, p = .010, \eta_p^2 = .198$]. The interaction between mask and delay was driven by the negative effect of the presence of a mask on hit rates irrespective of irrelevant feature randomization at the shortest delay period ($p = .003$, Bonferroni corrected). The irrelevant feature \times mask \times delay interaction was driven by a significant difference between mask and no mask hit rates in the randomized condition with a 1,300-ms delay ($p = .005$, Bonferroni corrected to alpha level .006).

Discussion

The purpose of Experiment 3 was to test whether preventing participants from relying on an uninterrupted, high-capacity memory trace to efficiently compare the sample and test displays abolishes the irrelevant change effect. To do this, visual pattern masks were briefly presented at the locations of each sample display item shortly after sample display offset. In line with our hypothesis, presenting a visual pattern mask shortly after the sample display removed the performance advantage observed in the no mask conditions. In keeping with the results of Experiments 1 and 2, this finding suggests that the dissipation of the irrelevant change effect over time likely does not reflect the gradual removal, either through suppression or passive decay, of task-irrelevant features from VWM representations.

The false-alarm and hit-rate analyses also revealed a pattern consistent with Experiments 1 and 2. False-alarm rates in the randomized condition tended to stay flat and hit rates did not increase over time. The mask did tend to produce a negative effect on hit rates for both the randomized and unchanged condition at the shortest delay. This could be because the short time period between the mask and response prevents adequate

consolidation into VWM and is more disruptive than randomization alone. Regardless, lower false-alarm rates were observed in the unchanged condition when no mask was present, but this effect disappeared when a mask was present, suggesting that the mask disrupts the same memory representation as the randomization.

General discussion

The current study tested the suppression explanation of the irrelevant change effect, which states that all of the features of attended objects are automatically bound together, but that over time task-irrelevant features can be removed from VWM representations through a process of top-down suppression. To test this hypothesis, in Experiments 1 and 2 participants were required to perform a demanding backward counting task concurrent with performance of a standard or irrelevant-change version of a change detection task requiring memory for bindings between color and location (Experiment 1) or shape and location (Experiment 2). If the elimination of the irrelevant change effect at longer delays reflects the use of executive resources to suppress task-irrelevant features, then occupying this resource should lead to a prolonged irrelevant change effect. Contrary to this proposal, in both experiments, taxing executive resources with a concurrent executive load task produced a large decrement in overall performance, but had no effect on the magnitude or time course of the irrelevant change effect. This finding suggests that executive resources are not needed for the effect to dissipate over time.

The analyses of hits and false alarms in the first two experiments revealed that the false-alarm rate increased over time when irrelevant features remained unchanged from sample to test, but there was no change in false alarms over time when irrelevant features were randomized. That is, when the task-irrelevant feature is randomized, the tendency to report a change when none has occurred remains stable over time, but when there is no randomization, this tendency increases over time. This suggests that the irrelevant change effect is likely caused by inflated performance in the unchanged condition at very short delays, rather than a reduction in the disruptive effect of irrelevant feature changes at longer delays. This finding is inconsistent with both top-down suppression as well as passive decay-based explanations of the irrelevant change effect. Instead, it suggests that participants likely made use of a high capacity, rapidly decaying form of visual memory to efficiently match the sample and test displays, a possibility that was confirmed by the results of Experiment 3. Specifically, when visual pattern masks were introduced between the sample and test displays, the irrelevant change effect disappeared at all time intervals. This supports the hypothesis that the irrelevant change effect stems from enhanced performance when a direct comparison is made between a fading

sensory representation and test displays in the unchanged condition. Interrupting this memory trace eliminates the irrelevant change effect even at the shortest delays.

Although the present results do not provide support for the top-down suppression hypothesis, they should not be taken as suggesting a lack of voluntary control over the contents of VWM. Indeed, several studies have suggested a high degree of control over what gets encoded and ultimately maintained in VWM. For example, it is possible to select a subset of objects in a memory display to encode into VWM, based on the surface features of the objects (e.g., encode blue and ignore red). This filtering ability is highly correlated with neural measures of the number of objects retained in VWM (Vogel & Machizawa, 2004). In addition to exerting voluntary control over which objects ultimately get encoded and maintained, other research suggests that it may be possible to preferentially store only task-relevant features of an object, depending on task goals (Serences, Ester, Vogel, & Awh, 2009; Woodman & Vogel, 2008) and the availability of low-level grouping cues (van Lamsweerde, Beck, & Johnson, 2016; Vogel et al., 2006). For example, van Lamsweerde, Beck, and Johnson (2016) showed that when detecting changes to displays of colored shapes, subjects could choose to encode either one or both feature dimensions, depending on whether a single dimension or both feature dimensions were task relevant. Additionally, studies of retro-cuing (i.e., cueing subjects to attend to a particular item from a memory display in the delay interval) show that available resources can be preferentially allocated to, and withdrawn from, VWM representations depending on their likelihood of being probed at test (Landman, Spekrijse, & Lamme, 2003; Lepsien, Griffin, Devlin, & Nobre, 2005; Lepsien & Nobre, 2006; Makovski & Jiang, 2007; Makovski, Sussman, & Jiang, 2008; Pertzov, Bays, Joseph, & Husain, 2013). Further evidence suggests that it may also be possible to remove task-irrelevant features from VWM representations once they have been formed. For example, Ye and colleagues (Ye, Hu, Ristaniemi, Gendron, & Liu, 2016) showed that a retro-cue directing participants to a single feature dimension increased the probability of correctly reporting a target from the cued dimension, but not the uncued dimension. The authors conclude that internal attention can be used to flexibly allocate available executive resources to specific feature dimensions within VWM representations. Thus, although the top-down suppression view does not adequately explain the elimination of the irrelevant change effect over time, this should not be taken to mean that voluntary control over the contents of VWM, both at the object and individual feature level, is not possible.

The irrelevant change effect has been of interest primarily because it is thought to provide evidence supporting object-based views of the nature of storage in VWM (Logie et al., 2011; Shen et al., 2013). Specifically, it has been argued that irrelevant feature changes are disruptive because task-

irrelevant features are encoded together with task-relevant features within an object-based VWM representation. The premise of this argument is that changes to task-irrelevant features could not exert a disruptive effect on performance if they were not remembered. However, the results presented here suggest that the irrelevant change effect does not reflect disruption due to irrelevant feature randomization, but a boost to performance due to highly accurate stimulus matching. The rapidly decaying, readily maskable nature of this benefit suggests that this boost likely relies on the use of a high capacity iconic or fragile VWM trace (Vandenbroucke & Sligte, 2011), rather than a more abstract, limited capacity, and stable VWM representation. Therefore, it is unclear from observing the irrelevant change effect whether the task-irrelevant features are stored in VWM, and, consequently, whether the VWM representation is feature- or object-based. However, the irrelevant change effect may still offer unique insights regarding the nature of visual memory. For example, the lack of evidence for time-based decay in the randomized displays could suggest that the decay generally attributed to VWM (Zhang & Luck, 2009) may partially reflect decay of this sensory representation. This possibility, however, requires additional research.

In this study, subjects monitored changes to feature bindings (e.g., color-location bindings). The issue of how feature bindings are remembered is also an area of considerable theoretical interest, specifically whether it is necessary to deploy attention to an object in VWM in order to maintain the feature bindings (Allen et al., 2006; Brown & Brockmole, 2010; Delvenne, Cleeremans, & Laloyaux, 2010; Fougne & Marois, 2009; Gajewski & Brockmole, 2006; Johnson, Hollingworth, & Luck, 2008; van Lamsweerde & Beck, 2012; Wheeler & Treisman, 2002). Although the current study was not designed to specifically address the role of attention in maintaining feature bindings, it might be predicted that, if attention is necessary, then a concurrent attentional load should eliminate the irrelevant change effect, as the bindings would be disrupted by the load task and the irrelevant changes should no longer have an effect on performance. Contrary to this possibility, the irrelevant change effect was present at the same time periods in both load conditions. This is broadly consistent with previous research showing that attention is not required to maintain feature bindings in working memory (Allen et al., 2006; Delvenne et al., 2010; Gajewski & Brockmole, 2006; Johnson et al., 2008; van Lamsweerde & Beck, 2012). On the other hand, concurrent performance of the high-load task did produce an overall disruption in memory for feature-location bindings at all study-test intervals examined. However, whether this effect is specific to memory for feature bindings or would apply equally to feature memory will require further research.

The present results support a novel hypothesis regarding the irrelevant change effect and its reduction over time.

Contrary to the suppression hypothesis, results showed that reducing the availability of executive resources did not have an impact on the duration of the irrelevant change effect, as would be expected if these resources were necessary to actively suppress the task-irrelevant features. Instead, our findings suggest that the source of the effect lies in the ability of participants to use a high-resolution, sensory memory representation to directly match the memory and test displays when no task-irrelevant feature changes are present.

Acknowledgements This research was supported by the National Institute of Health under Grant R15 MH105866-01 awarded to J.S.J.

Compliance with ethical standards

Conflicts of interest The authors declare no conflicts of interest.

References

- Allen, R. J., Baddeley, A. D., & Hitch, G. J. (2006). Is the binding of visual features in working memory resource-demanding? *Journal of Experimental Psychology: General*, 135, 298–313. doi:10.1037/0096-3445.135.2.298
- Baddeley, A. D. (2012). Working memory: Theories, models, and controversies. *Annual Review of Psychology*.
- Bradley, C., & Pearson, J. (2012). The sensory components of high-capacity iconic memory and visual working memory. *Frontiers in Psychology*, 3. doi:10.3389/fpsyg.2012.00355
- Brown, L., & Brockmole, J. R. (2010). The role of attention in binding visual features in working memory: Evidence from cognitive ageing. *The Quarterly Journal of Experimental Psychology*, 63(10), 2067–2079.
- Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional templates in visual working memory. *Journal of Neuroscience*, 31(25), 9315–9322. doi:10.1523/JNEUROSCI.1097-11.2011
- Conway, A. R. A., & Engle, R. W. (1994). Working memory and retrieval: A resource-dependent inhibition model. *Journal of Experimental Psychology: General*, 123(4), 354–373.
- Delvenne, J.-F. F., Cleeremans, A., & Laloyaux, C. C. (2010). Feature bindings are maintained in visual short-term memory without sustained focused attention. *Experimental Psychology*, 57(2), 108–116.
- Ecker, U. K. H., Maybery, M., & Zimmer, H. D. (2013). Binding of intrinsic and extrinsic features in working memory. *Journal of Experimental Psychology General*, 142(1), 218–234. doi:10.1037/a0028732
- Fougne, D., & Marois, R. R. (2009). Attentive tracking disrupts feature binding in visual working memory. *Visual Cognition*, 17(1-2), 48–66.
- Gao, Z., Li, J., Yin, J., & Shen, M. (2010). Dissociated mechanisms of extracting perceptual information into visual working memory. *PLoS ONE*, 5(12), 1–15. doi:10.1371/journal.pone.0014273
- Gajewski, D. A., & Brockmole, J. R. (2006). Feature bindings endure without attention: Evidence from an explicit recall task. *Psychonomic Bulletin & Review*, 13(4), 581–587.
- Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009). High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention. *Science*, 324(5931), 1207–1210.
- Hyun, J., Woodman, G. F., Vogel, E. K., Hollingworth, A., & Luck, S. J. (2009). The Comparison of Visual Working Memory

- Representations with Perceptual Inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35(4), 1140–1160. doi:10.1037/a0015019
- Jaswal, S. (2012). The importance of being relevant. *Cognitive Science*, 3, 309. doi:10.3389/fpsyg.2012.00309
- Jaswal, S., & Logie, R. H. (2011). Configural encoding in visual feature binding. *Journal of Cognitive Psychology*, 23(5), 586–603. doi:10.1080/20445911.2011.570256
- Johnson, J. S., Hollingworth, A., & Luck, S. J. (2008). The role of attention in the maintenance of feature bindings in visual short-term memory. *Journal of Experimental Psychology: Human Perception and Performance*, 34(1), 41–55
- Kane, M. J., & Engle, R. W. (2003). Working-memory capacity and the control of attention: The contributions of goal neglect, response competition, and task set to Stroop interference. *Journal of Experimental Psychology: General*, 132(1), 47–70. doi:10.1017/CBO9781107415324.004
- Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, 43(2), 149–164. doi:10.1016/S0042-6989(02)00402-9
- Lepsien, J., & Nobre, A. C. (2006). Cognitive control of attention in the human brain: Insights from orienting attention to mental representations. *Brain Research*, 1105(1), 20–31. doi:10.1016/j.brainres.2006.03.033
- Lepsien, J., Griffin, I. C., Devlin, J. T., & Nobre, A. C. (2005). Directing spatial attention in mental representations: Interactions between attentional orienting and working-memory load. *NeuroImage*, 26(3), 733–743. doi:10.1016/j.neuroimage.2005.02.026
- Logie, R. H., Gilhooly, K. J., & Wynn, V. (1994). Counting on working memory in arithmetic problem solving. *Memory & Cognition*, 22(4), 395–410. doi:10.3758/BF03200866
- Logie, R. H., Brockmole, J. R., & Jaswal, S. (2011). Feature binding in visual short-term memory is unaffected by task-irrelevant changes of location, shape, and color. *Memory & Cognition*, 39(1), 24–36.
- Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review*, 14(6), 1072–1078.
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(4), 369–380.
- Matsukura, M., & Hollingworth, A. (2011). Does visual short-term memory have a high-capacity stage? *Psychonomic Bulletin & Review*, 18(6), 1098–1104. doi:10.3758/s13423-011-0153-2
- Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology*, 39(5), 1224–1231.
- Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Attention, Perception, & Psychophysics*, 16, 283–290.
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, 20, 207–214.
- Shen, M., Tang, N., Wu, F., Shui, R., & Gao, Z. (2013). Robust object-based encoding in visual working memory. *Journal of Vision*, 13(2), 1. doi:10.1167/13.2.1
- Sligte, I. G. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in Psychology*, 1, 1–10.
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLoS One*, 3(2), e1699.
- Treisman, A., & Zhang, W. (2006). Location and binding in visual working memory. *Memory & Cognition*, 34(8), 1704–1719. doi:10.3758/BF03195932
- Unsworth, N., & Engle, R. W. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review*, 114(1), 104–132.
- van Lamsweerde, A. E., & Beck, M. R. (2012). Attention shifts or volatile representations: What causes binding deficits in visual working memory? *Visual Cognition*, 20(7), 771–792.
- van Lamsweerde, A. E., Beck, M. R., & Johnson, J. S. (2016). Visual working memory organization is subject to top-down control. *Psychonomic Bulletin & Review*, 23, 1181–1189.
- Vandenbroucke, A. R. E., & Sligte, I. G. (2011). Manipulations of attention dissociate fragile visual short-term memory from visual working memory. *Neuropsychologia*, 49, 1559–1568.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428, 748–751.
- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 32(6), 1436–1451.
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131(1), 48–64.
- Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin & Review*, 15(1), 223–229.
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2012). Flexibility in visual working memory: Accurate change detection in the face of irrelevant variations in position. *Visual Cognition*, 20(1), 1–28.
- Xu, Y. (2010). The neural fate of task-irrelevant features in object-based processing. *Journal of Neuroscience*, 30(42), 14020–14028. doi:10.1523/JNEUROSCI.3011-10.2010
- Ye, C., Hu, Z., Ristaniemi, T., Gendron, M., & Liu, Q. (2016). Retro-dimension-cue benefit in visual working memory. *Scientific Reports*, 6(September), 35573. doi:10.1038/srep35573
- Yin, J., Gao, Z., Jin, X., Ye, L., Shen, M., & Shui, R. (2011). Tracking the mismatch information in visual short term memory: An event-related potential study. *Neuroscience Letters*, 491(1), 26–30. doi:10.1016/j.neulet.2011.01.001
- Yin, J., Gao, Z., Jin, X., Ding, X., Liang, J., & Shen, M. (2012). The neural mechanisms of percept–memory comparison in visual working memory. *Biological Psychology*, 90(1), 71–79. doi:10.1016/j.biopsycho.2012.02.023
- Yin, J., Zhou, J., Xu, H., Liang, J., Gao, Z., & Shen, M. (2012). Does high memory load kick the task irrelevant information out of visual working memory? *Psychonomic Bulletin & Review*, 19(2), 218–224.
- Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, 20, 423–428.
- Zhou, J., Yin, J., Chen, T., Ding, X., Gao, Z., & Shen, M. (2011). Visual Working Memory Capacity Does Not Modulate the Feature-Based Information Filtering in Visual Working Memory. *PLoS ONE*, 6(9). doi:10.1371/journal.pone.0023873